LLNL-PROC-644577

The origin of mass. Update, October 2013.

P. Boyle, M. Buchoff, N. Christ, T. Izubuchi, C. Jung, Z. Lin, T. Luu, R. Mawhinney, C. Schroeder, R. Soltz, P. Vranas, J. Wasem, H. Yin

October 4, 2013

Supercomputing 2013
Denver, CO, United States
November 17, 2013 through November 22, 2013

Disclaimer

The origin of mass. Update, October 2013.

Peter Boyle*, Michael I. Buchoff†, Norman Christ‡, Taku Izubuchi§, Chulwoo Jung§, Zhongjie Lin‡, Thomas C. Luu†, Robert Mawhinney‡, Chris Schroeder†, Ron Soltz†, Pavlos Vranas†, Joseph Wasem†, Hantao Yin‡
*SUPA, School of Physics, The University of Edinburgh, Edinburgh EH9 3JZ, UK
†Physical Sciences Directorate, Lawrence Livermore National Laboratory, Livermore, California 94550, USA
‡Physics Department, Columbia University, New York, NY 10027, USA
§Physics Department, Brookhaven National Laboratory, Upton, NY 11973, USA

***Abstract*—Since our April Gordon Bell Prize submission we have proceeded to the exploration of the order of the QCD phase transition as the constituent quark masses are varied to even lower values. This is significant to the understanding of QCD as well as to a possible cosmic transition that could have created composite dark matter. Here we report a time-to-solution over the previous-state-of-the-art that is approximately a factor of 2 better than our April submission and was made possible because of our algorithms. This is approximately 400 times better than the previous state of the art. This opens a new area of research that was previously out of reach. Here we provide first physics results from simulations on the LLNL BlueGene/Q systems. We also show that the factor of 400 is robust and obtainable on the full LLNL BlueGene/Q system.**

Category: **Scalability and Time to solution**

I. Background

The phase diagram of the QCD thermal transition as a function of the mass of the up and down quark masses (set to be equal here since they are nearly equal in nature; x axis) and the strange quark mass (y axis) is shown in Figure 1. The diamond indicates the physical point where these quarks have their experimentally measured values. At that point the transition is a crossover from the quark gluon plasma to the stable nuclear matter of today. This diagram is referred as the Columbia phase diagram [1]. The remaining three quarks(charm, top, and bottom) are considerably heavier and do not affect the order of the transition. If dark matter is composite and governed by a similar theory as QCD then its cosmic thermal transition would correspond to a point in a similar plot.

The upper right part of the phase diagram corresponding to large quark masses is well established by early numerical simulations of Lattice QCD. This is easy to simulate because the heavy quark masses directly set the lowest eigenvalues of the Dirac matrix to large values making its inversion very fast. The bottom right and top left are similarly well understood.

The bottom left of this diagram is the most important physically. Not only does it contain the physical point (diamond) but also is expected to have a richer phase structure consisting of second order phase transition lines, a tricritical point and a first order phase transition region. That phase structure is important since it affects the more detailed properties of the nearby physical point (diamond) as well as it indicates the sensitivity on the quark masses.

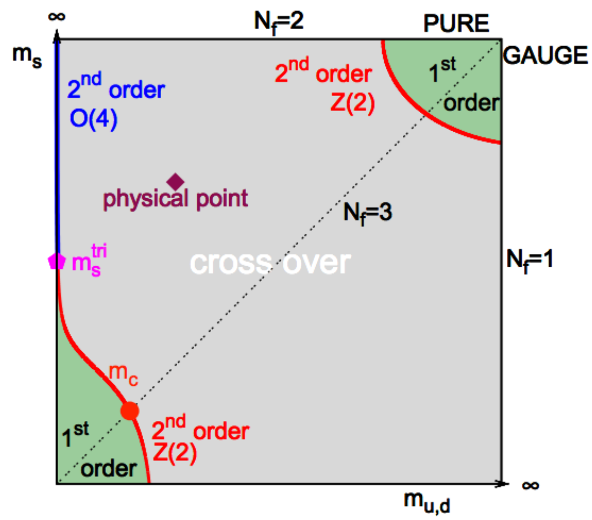

Fig. 1. The phase diagram of the QCD thermal transition as a function of the quark masses, Columbia phase diagram [1], from [2].

The problem is that although we now know that the physical point is in the crossover region the rest of the bottom left is largely a theoretical conjecture. In addition, the lines shown are schematic. Even if the conjecture proves to be right we do not know to which values of the quark masses they correspond. For example how close is the physical point to the second order line? Or how small should the quark masses be to enter the first order transition region? These questions can only be answered using Lattice QCD simulations. But this has been impossible over the past nearly 40 years because the corresponding quark masses are small demanding very large lattices and having very ill-conditioned matrices that have been prohibitive to invert. In addition that region demands firm control over chiral symmetry, a fundamental ingredient that ensures the pions have equal masses. This is spoiled by the lattice and could not be faithfully addressed until the advent of Domain Wall Fermions that add additional computational costs. It is this problem that for the first time we are able to investigate. Because of our time-to-solution speedup we estimate that on LLNL BG/Q Sequoia type resources it would take about half a year to complete this project. Just a few years ago this project would have taken 200 years to complete.

II. Time-to-solution

Since our original submission in April we have been applying our algorithms and code implementation on the BlueGene/Q LLNL supercomputers to this problem. As a first step we were able to decrease the pion masses to $m_\pi = 100$ MeV in order to get closer to the conjectured second order line (in our April submission the pion mass was set at its physical value $m_\pi = 140$ MeV). This was possible because of the properties of our algorithms (see section 4.1 of our April submission) and was achieved by exploring their parameter space and finding the necessary values for the Mobius parameter and input quark mass. In particular, as explained there, our DSDR algorithm suppresses dislocations that are lattice artifacts. This suppression allows us to further adjust the parameters of our algorithm and achieve this low pion mass. As described below this decreases our time to solution by another factor of 2 for a total of 400 times speedup over the previous state of the art. The value of $m_\pi = 100$ MeV is estimated using standard DWF methods.

The computational cost of the previous state-of-the-art for this problem can be estimated from the results in [3]. There it was found that for regular DWF, in order to obtain $m_\pi = 225$ MeV at temperature $T \approx 140$ MeV a fifth dimension with $L_s = 96$ sites is needed. In addition there it was shown that at these parameters $m_\pi^2 \sim \frac{1}{L_s}$. Therefore in order to achieve $m_\pi = 100$ MeV one would need $L_s \approx 96(\frac{225}{100})^2 = 486$. In contrast, using our algorithms as described in our original submission after adjusting the Mobius parameter (from $c = 1.5$ to $c = 2.0$) and the quark mass (from $m = 0.002$ to $m = 0.001$) we were able to reach $m_\pi = 100$ MeV with $L_s = 24$ at this temperature. Because computational cost is directly proportional to L_s this is a speedup of approximately 20 over the previous state of the art. This is twice as large as our factor of 10 in our April submission. Our code efficiency for this problem has remained the same as in April and it is a factor of 20 over the previous 1 Petaflops systems as described in our April submission. These two improvements result to a factor of 400 over the previous state-of-the-art.

In April the largest accessible machine at LLNL was the Sequoia+Vulcan 120 rack BG/Q supercomputer. Our April submission was done based on that system. Since April Sequoia and Vulcan were separated and Sequoia became classified and non-accessible. The simulation results presented here are therefore from the 24 rack Vulcan system and in particular from an 8-rack dedicated partition that LLNL has provided us. However, our time-to-solution result is robust. If we had access to the full original 120 rack system then we could have achieved the same time-to-solution scaled by $8/120$. The exploration of the phase diagram has inherently three extra dimensions: the ranges of the quark masses (as in figure 1) and the temperature range needed to determine the order of the transition. Because our implementation has strong scaling that is linear up to 16 racks the most efficient implementation of our project on 120 racks would involve 5 different quark mass values at 3 different values of the temperature simulated at the same time on 15 8-rack systems. Of course, we could also use the full 120 rack system for a single simulation on two 60 rack systems (one with initial Molecular Dynamics state in each of the two possible phases) and incur the strong scaling degradation of figure 7 in our April submission which is about 30%. Therefore our science problem is robust and can adapt to system size with nearly no losses.

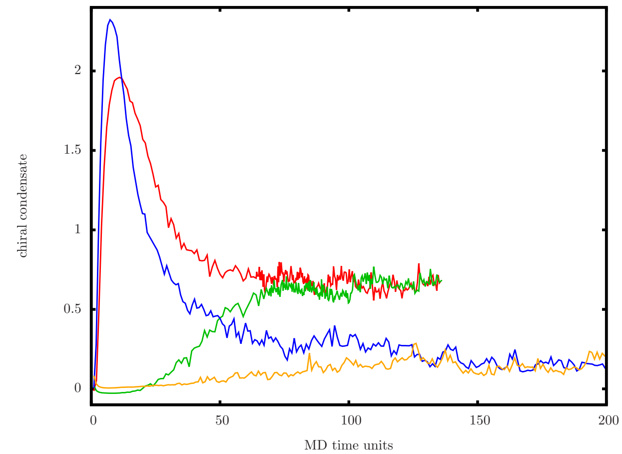

Fig. 2. The evolution of the order parameter (chiral condensate) as a function of the molecular dynamics evolution time units for pion mass of 100 MeV at $T = 140$ MeV (top two lines, red and green) and at $T = 160$ MeV (bottom two lines, blue and orange). For each temperature the lines that start lower (green and orange) correspond to an original MD state in the high temperature phase while the lines that start higher (red and blue) correspond to an original MD state at the zero temperature phase. The lattice size is $64^3 \times 8$.

III. First physics results

In this section we present our first physics results from simulations at $T = 140$ MeV and $T = 160$ MeV with $m_\pi = 100$ MeV and physical strange quark mass on large $64^3 \times 8$ lattices. All simulation where done on an 8 rack BG/Q system at LLNL. This corresponds to a horizontal departure to the left of the physical point (diamond in figure 1). These results demonstrate that this research is now a reality.

The order of a phase transition is determined by the behavior of the order parameter. The order parameter changes abruptly as the temperature is varied past the transition. By calculating the way it changes we can determine the location and order of the transition. An example of a familiar order transition is the phase change of water to steam as the temperature increases. An order parameter for that transition would be the number of water molecules per unit volume.

For the QCD thermal transition the order parameter is the chiral condensate. For a first order transition the order parameter should evolve differently depending on the original state of the molecular dynamics (MD) evolution. Even after a large number of MD iterations the system would relax at different values of the order parameter depending on the original MD state. This two state signal typically persist for large MD times and can abruptly tunnel from one phase to the other. This behavior would be a clear indication of a first order transition.

Our first results at temperatures $T = 140$ MeV and $T = 160$ MeV are shown in Figure 2. Notice that the chiral condensate for $T = 140$ MeV relaxes at a value that is about three times larger than that at $T = 160$ MeV. This rapid change indicates that the transition is occurring between these temperatures. Clearly there is no sign of a two state signal in ether temperature. For each temperature the system relaxes to the same value independently of the original state. This indicates that the transition for $m_\pi = 100$ MeV may still be in the crossover region. However, this is not yet conclusive. The order of the transition is determined at the transition point. Therefore we need to locate the temperature where the transition occurs (likely between $T = 140$ MeV and $T = 160$ MeV) and observe the MD evolution of the order parameter at that temperature. This research is currently underway on LLNL's BG/Q Vulcan machine.

Acknowledgments

The work of MB, TL, CS, RS, PV, and JW was performed under the auspices of the U.S. Department of Energy by LLNL under Contract No. DE-AC52-07NA27344. This research was partially supported by the LLNL LDRD "Unlocking the Universe with High Performance Computing" 10-ERD-033 and by the LLNL LDRD "Illuminating the Dark Universe with Petaflops Supercomputing" 13-ERD-023. We warmly thank B. Carnes and the Computation Directorate at LLNL for providing access to the Vulcan LLNL supercomputer.

References

[1] Brown et. al., "On the existence of a phase transition for QCD with three light quarks", *Phys. Rev. Lett.* 65, 2491 (1990)

[2] H.-T. Ding et. al., "Exploring phase diagram of $N_f = 3$ QCD at $\mu = 0$ with HISQ fermions", *PoS LATTICE2011* 191 (2011), arXiv:1111.0185 [hep-lat]

[3] A. Bazavov et.al, "The chiral transition and $U(1)_A$ symmetry restoration from lattice QCD using Domain Wall Fermions", *Phys. Rev.* D86, 094503 (2012).